ARMY RESEARCH LABORATORY

Scalability of the CTH Shock Physics Code on the Cray XT

by Stephen J. Schraml and Thomas M. Kendall

ARL-RP-281   **November 2009**

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

Army Research Laboratory
Aberdeen Proving Ground, MD 21005-5066

ARL-RP-281   November 2009

Scalability of the CTH Shock Physics Code on the Cray XT

Stephen J. Schraml and Thomas M. Kendall
Weapons and Materials Research Directorate, ARL

REPORT DOCUMENTATION PAGE

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
November 2009	Reprint	September 2008–February 2009

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Scalability of the CTH Shock Physics Code on the Cray XT	
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)	5d. PROJECT NUMBER
Stephen J. Schraml and Thomas M. Kendall	1L162618AH80
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
U.S. Army Research Laboratory ATTN: RDRL-WMT-C Aberdeen Proving Ground, MD 21005-5066	ARL-RP-281

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

A reprint from the *2009 DOD High Performance Computing Users Group Conference,* San Diego, CA, 15–18 June 2009.

14. ABSTRACT

This report presents an overview of an explicit message-passing paradigm for a Eulerian finite-volume method for modeling solid dynamics problems involving shock wave propagation, multiple materials, and large deformations. Three-dimensional simulations of high-velocity impact were conducted on two scalable, high-performance computing systems to evaluate the performance of the message-passing code. Simulations were performed using >3 billion computational cells running on more than 8000 processor cores. The performance of the message-passing code was found to scale linearly on the two computer systems evaluated across the range of cases considered.

15. SUBJECT TERMS

penetration, computer simulation, finite-element method, shock waves, impact, numerical simulation, high-performance computing

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE	UL	14	Stephen J. Schraml
Unclassified	Unclassified	Unclassified			19b. TELEPHONE NUMBER *(Include area code)* 410-278-6556

Scalability of the CTH Shock Physics Code on the Cray XT

Stephen J. Schraml, stephen@arl.army.mil
Thomas M. Kendall, tkendall@arl.army.mil
U.S. Army Research Laboratory
Aberdeen Proving Ground, MD 21005-5066

1   Introduction

During the 2007 and 2008 Technology Insertion initiatives (TI07 & TI08), the Department of Defense (DoD) High Performance Computing Modernization Program (HPCMP) acquired and deployed several High Performance Computing (HPC) systems in the Cray XT series. Two of these systems are the Cray XT4 located at the U.S. Army Engineer Research and Development Center (ERDC) DoD Supercomputing Resource Center (DSRC) and the Cray XT5 located at the U.S. Army Research Laboratory (ARL) DSRC. Each DSRC sponsored a pioneer user access initiative for the purpose of conducting performance studies of the two systems prior to production utilization. During this pioneer user period, the scalable performance of the CTH shock physics code was evaluated on each system.

CTH is an explicit, Eulerian, finite volume code for the numerical simulation of the high-rate response of materials to impulsive loads [1]. CTH is widely used across the defense research and development complex in the development of explosives, blast and fragmenting warheads, kinetic energy penetrators, vehicle armor systems, protective structures, etc. CTH employs a single program multiple data (SPMD) programming model to employ scalable performance on distributed memory HPC systems. CTH has been found to achieve linear scalability in performance on many systems deployed by the HPCMP since the late 1990s [2, 3, 4, 5, 6].

2   HPC Systems

The Cray XT4 system installed at the ERDC DSRC is configured with 2152 compute nodes, each of which has a 2.1-GHz Advanced Micro Devices (AMD) quad-core processor, resulting in a total of 8608 processor cores. Each compute node has 8 GB of memory, resulting in average memory availability of 2 GB per processor core, and a total system memory of approximately 16.4 TB [7]. When the XT4 was initially installed, it was configured with dual-core AMD processors. The compute nodes were later upgraded to quad-core processors. The CTH scalability study was performed on the XT4 for both dual-core and quad-core configurations.

The Cray XT5 system installed at the ARL DSRC is configured with 1300 compute nodes, each of which has two 2.3-GHz AMD quad-core processors, resulting in a total of 10,400 processor cores. Each compute node has 32 GB of memory, resulting in an average memory availability of 4 GB per processor core, and a total system memory of 41.6 TB. Both the XT4 and XT5 employ a proprietary high-speed network for communication between compute nodes [8].

3 Scalability Study Parameters

The scalability of CTH on the XT systems was determined through a series of simulations that employed both fixed and adaptive meshes. The benchmark simulation that was employed for the study involved the yawed, oblique impact of a depleted uranium long rod penetrator against a steel target plate. This benchmark problem was selected because of its relevance to DoD problems in shock physics. A thorough description of the benchmark simulation has been previously documented [2, 3, 4, 5, 6].

The fixed-mesh scalability simulations were conducted with a nearly constant workload. This was done to keep the computation-to-communication ratio as close to constant as possible for simulations involving different numbers of processor cores (one CTH task assigned to each processor core). Maintaining a nearly constant computation-to-communication ratio and minimizing disk access for intermediate plot and restart files during the time integration permitted the computational performance to be isolated and measured as a function of the number of processor cores used.

As the number of CTH tasks was increased, the fixed mesh was incrementally refined by uniformly decreasing the characteristic cell size in each coordinate direction by the nearest integer factor of $2^{-1/3}$. This approach approximately doubles the total number of Eulerian cells with each successive mesh refinement. The characteristics of the meshes used in the scalability study are summarized in Table 1. In this table, the columns NI, NJ, and NK refer to the number of Eulerian cells in the x, y, and z directions, respectively. The mesh sizes listed in the table produce computational sub-domains containing approximately 387,000 Eulerian cells each. For the 8192-task simulation, this results in a computational domain containing approximately 3 billion Eulerian cells.

The adaptive mesh refinement (AMR) capability in CTH allows the definition of the mesh to change during the simulation based on the evolving characteristics of the simulation [9]. The adaptation of the mesh is based on user-defined indicators, such as the value, gradient, or difference, of a variable in the solution (pressure, density, velocity, stress, etc.). This technique results in simulations in which the most highly resolved mesh "follows" the activity of interest to the analyst while using less highly resolved mesh in the remainder of the computational domain. This

Table 1. Scalability study parameters.

Number of Tasks	Fixed Mesh					Adaptive Mesh	
	NI	NJ	NK	Total Number of Cells	Cell Size (mm)	Maximum Refinement Level	Minimum Cell Size (mm)
1	215	30	60	387,000	1.000	5	0.750
2	271	38	75	772,350	0.793	5	0.750
4	341	48	95	1,554,960	0.630	5	0.750
8	430	60	120	3,096,000	0.500	6	0.375
16	541	76	151	6,208,516	0.397	6	0.375
32	683	95	191	12,393,035	0.315	6	0.375
64	860	120	240	24,768,000	0.250	7	0.188
128	1083	151	302	49,386,966	0.199	7	0.188
256	1366	190	382	99,144,280	0.157	7	0.188
512	1720	240	480	198,144,000	0.125	8	0.094
1024	2166	302	604	395,095,728	0.099	8	0.094
2048	2732	380	764	793,154,240	0.079	8	0.094
4096	3440	480	960	1,585,152,000	0.063	9	0.047
8192	4334	605	1210	3,172,704,700	0.050	9	0.047

allows the analyst to configure highly resolved simulations that have fewer total computational cells than a comparable fixed-mesh simulation having the same minimum cell size.

The AMR implementation in CTH is a block-based scheme in which each block consists of an orthogonal mesh with a fixed number of cells in the x, y, and z directions. The blocks are connected in a hierarchical manner with adjacent blocks having either exactly the same cell size or exactly a 2:1 ratio in cell size. Refinement or un-refinement of the mesh is accomplished through a series of transitions of adjacent blocks with a difference in mesh density of 2:1. All mesh blocks at a given mesh density are at the same refinement level. The finest mesh resolution that can exist in the computational domain is controlled by defining the maximum refinement level of the mesh.

The AMR CTH benchmark used in the scalability study was configured to be physically identical to the fixed-mesh simulation. The only difference between the fixed-mesh simulation and the AMR simulation was the definition of the mesh. The size of the mesh in the AMR simulation was scaled with the number of CTH tasks in a manner similar to the fixed-mesh study. However, it is not possible to precisely scale the total number of cells in the AMR simulation since the refinement and un-refinement indicators are based on the physics, not the topology of the computational domain. Thus, to scale the size of the simulation in a controlled manner, the maximum refinement level was increased by one for every factor of eight increase in the number of tasks. The 2:1 ratio of cell size between refinement levels results in a factor of approximately eight in the total number

of cells in the 3-D simulation. The variation of the maximum refinement level and the resulting minimum cell size with the number of tasks used is summarized in Table 1.

The scalable performance of CTH is measured by the "grind time," which is the average time required for the code to update all field variables for one computational cell in a given time increment (cycle). In a case of ideal scalability, the grind time will decrease by a factor of two for every doubling of number of processor cores used if the ratio of computation to communication is held constant.

4   Study Results: Cray XT4

The results of the CTH scalability study on the XT4 are provided in Figure 1. In this figure, the fixed mesh study results are plotted in the chart on the left of the figure and the adaptive mesh results are plotted in the chart on the right. Both dual-core and quad-core results are provided in each chart. The results demonstrate that CTH achieves linear scalability on the XT4 up to 2048 CTH tasks, the maximum number considered in the study on this system.
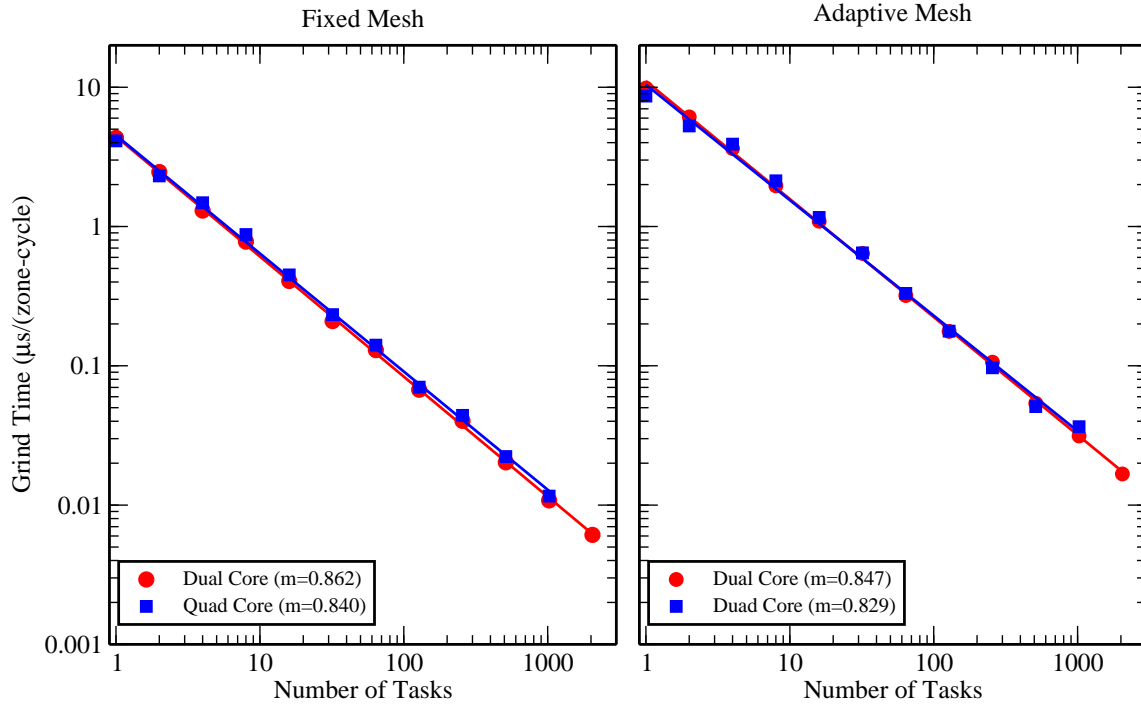
Figure 1. Scalability on XT4 with dual- and quad-core processors.

The CTH simulations run on the XT4 were executed in such a way as to completely fill the assigned compute nodes whenever possible. For example, a job using eight CTH tasks under the quad-core configuration would use two nodes.

The linear scalability of CTH can be described by $g_n = g_1/n^m$ where g_n is the expected grind time on n processor cores, g_1 is the measured grind time on one processor core, and m is the parallel efficiency, the slope of the scalability line. A slope of 1 corresponds to perfect scalability. A regression analysis was performed on the measured performance data to determine the parallel efficiency of CTH on the systems under consideration. In both fixed and adaptive mesh results, the parallel efficiency on the XT4 was slightly better for the dual-core processors than the quad-core processors. However, this performance penalty is far outweighed by the doubling in total system capability by replacing the dual-core processors with quad-core processors.

5 Study Results: Cray XT5

When the scalability simulations were run on the XT5, they were configured to use every possible power-of-two combination of nodes and tasks per node to achieve the desired total number of CTH tasks. For example, for the case involving 4 total CTH tasks, simulations were performed on: (1) one node using four processor cores, (2) two nodes using two processor cores per node, and (3) four nodes using one processor core per node. Executing the study in this way makes it possible to determine the effect of the number of CTH tasks assigned per node on the overall computational performance.

The results of the fixed mesh CTH scalability study on the XT5 are provided in Figure 2. The chart on the left of this figure is a plot of the grind time as a function of CTH tasks, similar to that provided in Figure 1. In this chart, the data are plotted separately for different numbers of CTH tasks assigned to each node. The plot shows that for a given number of tasks there is a slight increase in the grind time as the number of tasks per node is increased. This is likely a result of increased memory contention on the node. This slight penalty associated with increased tasks per node is counterbalanced by the fact that the parallel efficiency, m, increases slightly as the number of tasks per core is increased.

The chart on the right of Figure 2 attempts to quantify the effect of the number of tasks per node on the performance. In this plot, for each total number of CTH tasks assigned, from eight to 1024, the grind time is plotted relative to the one task-per-node case. This plot shows that as the number of tasks per node is increased to the maximum of eight tasks per node, the grind time increases by approximately 20%. However, the cases involving the largest processor counts (512 and 1024) had the lowest penalty on the grind time as the number of tasks per node was increased. Thus, for the
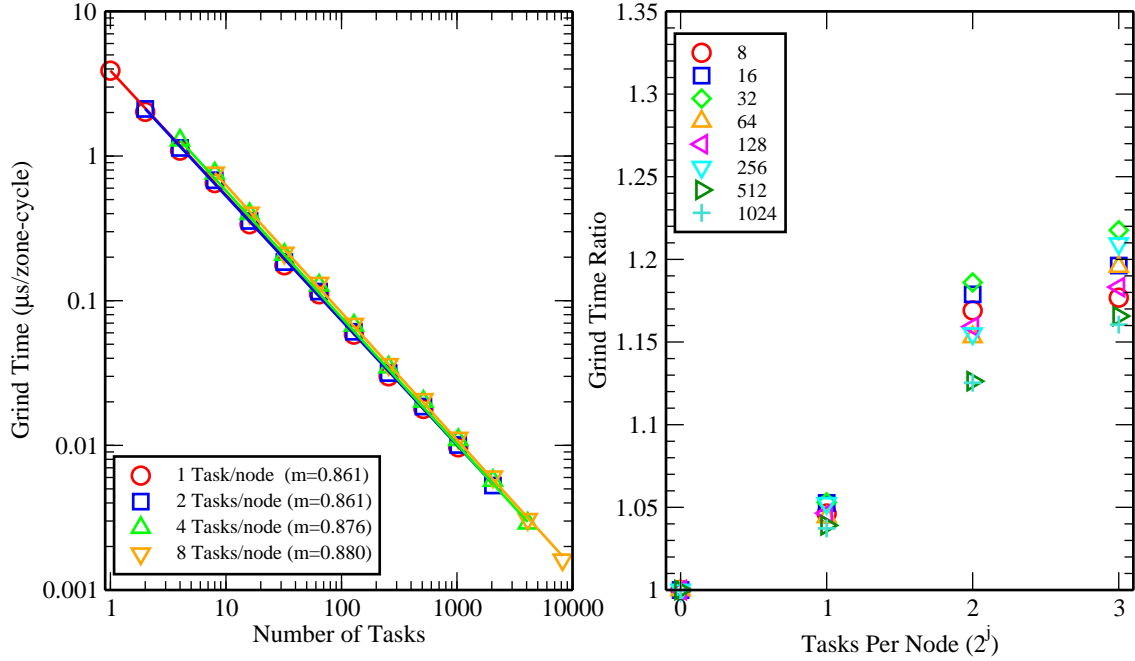
Figure 2. Fixed mesh scalability and relative performance on XT5.

typical operating environment in which all processor cores in the compute nodes are assigned a task, the XT5 system operates most efficiently when running large jobs.

The results of the adaptive mesh CTH scalability study on the XT5 are provided in Figure 3. This figure plots the results in the same format and at the same scale as the fixed mesh results in Figure 2. The findings of the adaptive mesh scalability study on the XT5 are similar to the fixed mesh findings: (1) the parallel efficiency generally increases as the number of tasks per node is increased, (2) the grind time ratio increases as the number of tasks per node is increased, indicating increased memory contention on the node, and (3) the grind time ratios are smallest for the largest number of tasks assigned.

6 Summary

Linear scalability of CTH on the Cray XT4 and XT5 systems has been demonstrated for simulations using up to 8192 processor cores. The linear scalability was demonstrated for simulations using both fixed and adaptive meshes. At the time of installation of these two systems, they were among the largest deployed under the auspices of the HPCMP. This fact, combined with the finding that these systems operate most efficiently for simulations that use large numbers of processor cores
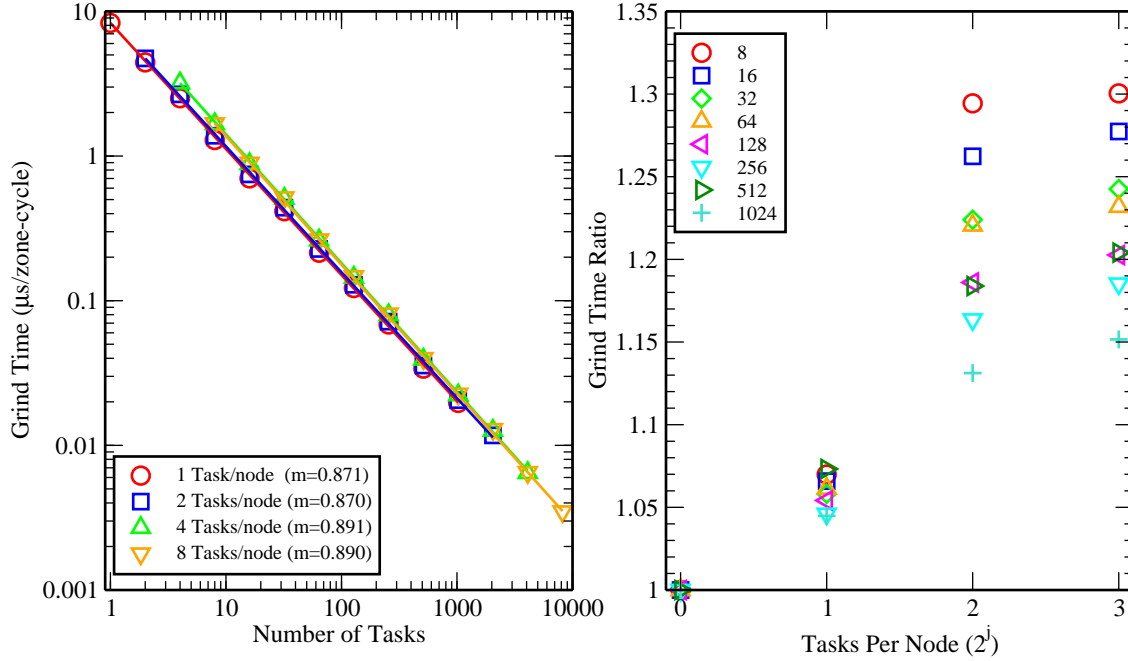
Figure 3. Adaptive mesh scalability and relative performance on XT5.

indicates that the productivity of these systems would be maximized by targeting large jobs to these resources.

7 Acknowledgment

References

1. McGlaun, J.M. and S.L. Thompson. "CTH: A Three-Dimensional Shock Wave Physics Code." *International Journal of Impact Engineering*, 1990, **10**, 351-360.

2. Kimsey, K.D., S.J. Schraml, and E.S. Hertel. "Scalable Computations in Penetration Mechanics." *International Journal on Advances in Engineering Software including Computing Systems in Engineering*, 1998, **29**, 209-215.

3. Schraml, S.J. and K.D. Kimsey. "Scalability of the CTH Hydrodynamics Code on the Sun HPC 1000 Architecture." U.S. Army Research Laboratory, ARL-TR-2173, February 2000.

4. Schraml, S.J., K.D. Kimsey, and T.M. Kendall. "Scalable Simulations of Penetration Mechanics on the SGI Origin 3800 and the IBM SP Power3 Computer Systems." U.S. Army Research Laboratory, ARL-TR-2537, June 2001.

5. Kendall, T.M., C.M. McCraney, and S.J. Schraml. "Early Experience With the IBM Power4 Scalable Parallel System at the ARL MSRC." Proceedings of the 2002 DoD High Performance Computing User's Group Conference, Austin, TX, June 2002.

6. Kendall, T.M. and S.J. Schraml. "Performance on CTH on a Myranet/Xeon Cluster." Cluster-World 2004, San Jose, CA, April 2004.

7. *Cray XT4 User's Guide.* www.erdc.hpc.mil

8. *Cray XT Series Application Programming and Optimization.* TR-XTPO-J-ARL, Cray Inc. September 2008.

9. Bell, R.L., R.M. Baer, R.M. Brannon, D.A. Crawford, M.G. Elrick, E.S. Hertel, R.G. Schmitt, S.A. Silling, and P.A. Taylor. "CTH User's Manual and Input Instructions, Version 7.0." Sandia National Laboratories, Albuquerque, NM, April 2005.

NO. OF
COPIES ORGANIZATION

1   DEFENSE TECHNICAL
(PDF   INFORMATION CTR
only)   DTIC OCA
        8725 JOHN J KINGMAN RD
        STE 0944
        FORT BELVOIR VA 22060-6218

1   DIRECTOR
        US ARMY RESEARCH LAB
        IMNE ALC HRR
        2800 POWDER MILL RD
        ADELPHI MD 20783-1197

1   DIRECTOR
        US ARMY RESEARCH LAB
        RDRL CIM L
        2800 POWDER MILL RD
        ADELPHI MD 20783-1197

1   DIRECTOR
        US ARMY RESEARCH LAB
        RDRL CIM P
        2800 POWDER MILL RD
        ADELPHI MD 20783-1197

        ABERDEEN PROVING GROUND

1   DIR USARL
        RDRL CIM G (BLDG 4600)

NO. OF
COPIES ORGANIZATION

  2   US ARMY ARDEC
        AMSRD AAR MEE W
        D PHAU
        C CHIN
        BLDG 3022
        PICATINNY NJ 07806 5000

  3   NSWC
        M HOPSON
        W CHEPREN
        C DYKA
        6138 NORC AVE STE 313
        DAHLGREN VA 22448 5157

  1   US ARMY ENGINEER RESEARCH
        AND DEVELOPMENT CENTER
        CEERD GM I
        S AKERS
        3909 HALLS FERRY RD
        VICKSBURG MS 39180-6199

        ABERDEEN PROVING GROUND

 11   DIR USARL
        RDRL CIH C
          J CAZAMIAS
        RDRL CIH S
          T KENDALL
        RDRL WMT A
          D DONEY
          D KLEPONIS
        RDRL WMT B
          R BANTON
          S KUKUCK
        RDRL WMT C
          T BJERKE
          K KIMSEY
          S SCHRAML
        RDRL WMT D
          S BILYK
          B LOVE